

RESEARCH

Open Access



A unified message-passing algorithm for MIMO-SDMA in software-defined radio

Alexander Kocian^{1*} , Mihai-Alin Badiu², Bernard Henri Fleury², Francesca Martelli³ and Paolo Santi^{3,4}

Abstract

This paper presents a novel software radio implementation for joint channel estimation, data decoding, and noise variance estimation in multiple-input multiple-output (MIMO) space division multiple access (SDMA). In contrast to many other iterative solutions, the proposed receiver is derived within the theoretical framework of a unified message-passing algorithm, combining belief propagation (BP) and the mean field approximation (MF) on the corresponding factor graph. The algorithm minimizes the region-based variational free energy in the system under appropriate conditions and, hence, converges to a fixpoint. As a use-case, we consider the high-rate packet-oriented IEEE 802.11n standard. Our receiver is implemented on a software-defined radio platform dubbed MIMONet, composed of a GNU radio software component and a universal software radio peripheral (USRP). The receiver was evaluated in real indoor environments. The results of our study clearly show that, once synchronization issues are properly addressed, the BP-MF receiver provides a substantial performance improvement compared to a conventional receiver also in real-world settings. Such improvement comes at the expense of an increase in running time that can be as high as 87. Therefore, the trade-off between communication performance and receiver complexity should be carefully evaluated in practical settings.

Keywords: MIMO communications, Space division multiple access, Belief propagation, Mean field approximation, Factor graphs, Software-defined radios

1 Introduction

Multiple-input multiple-output (MIMO) technology is popular in wireless communications due to the increased spectrum efficiency brought along by the use of multiple antennas in transmission, reception, or both. A further performance improvement is possible when MIMO technology is used in combination with orthogonal frequency division multiplexing (OFDM) modulation, namely, when different streams of information bits are modulated on orthogonal subcarriers.

The basic role of the receiver is to decode the information bits from the received signal which is affected by various unknown factors, such as the channel response and receiver noise. Receivers were originally designed to process the received signal in a cascaded fashion, starting with synchronization, then channel estimation, equalization, and finally decoding. Building on the intuition that

(soft) information generated by one module can actually be re-used as a refined input to other modules, a multitude of receiver structures performing iterative, “turbo”-like processing have been proposed in the literature. Results unequivocally show substantial performance gains compared to non-iterative architectures. However, separate design of the individual modules cannot provide the guarantee of global optimality or convergence. Moreover, it is not clear what type of information the modules should exchange and how to combine/process it. In recent years, several works have looked at the receiver design from the perspective of Bayesian inference on graphical models [1]. The use of formal frameworks for approximate inference allows for a principled design of iterative receiver structures. Among the various algorithms for approximate inference [1], the belief propagation (BP) algorithm [2] (also referred to as sum-product algorithm [3]) is the most celebrated one. The algorithm operates by passing messages on the so-called factor graph, which represents the factorization of the probabilistic model of the communication system. The algorithm is able to compute

*Correspondence: kocian@di.unipi.it

¹Department of Computer Science, University of Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy

Full list of author information is available at the end of the article

exact marginal probability density functions (pdfs) only when the factor graph is cycle free; otherwise, it outputs approximations of the marginal pdfs. Still, it was shown to work well in many graphs with cycles. This led to its widespread use in digital communications [3, 4]. However, BP usually yields intractable computations in probabilistic models with mixed discrete and continuous variables (see, e.g., [5]), which is also the case with our MIMO model. Different from BP, the mean field (MF) approximation [1] is a variational inference method that has been successfully applied to continuous probabilistic models. The algorithm also admits a message-passing formulation [6], and typically has simple closed-form message computation, especially for conjugate-exponential models. Its main drawback is that, for some models, the provided solution is not accurate enough, due to the underlying approximation of the joint posterior pdf (see [1] for more details). When the probabilistic model contains both continuous and discrete variables and the dependencies between them are both deterministic and stochastic, it is advantageous to apply the BP and MF algorithms in those parts of the factor graph where they are most suitable. For this purpose, we employ the recently proposed unifying inference framework that combines BP and MF [7]. In a nutshell, the factor graph is divided into two parts, the BP part and the MF part, where, respectively, BP-like and MF-like messages are computed. The framework states clear message computation rules, including specific expressions for the messages to be passed between the two parts. The unified nature of the algorithm resides in the fact that it iteratively optimizes a single objective function. Having derived a BP-MF-based MIMO-space division multiple access (SDMA) receiver in [8], we want to adapt and test it in real channels using software-defined radio (SDR).

SDR systems are becoming commonly used in the wireless networking research community due to their flexibility in rearranging different communication architectures with limited effort. Three key virtues of SDR are *reconfigurability*, *intelligence*, and *flexibility* [9]. To reduce hardware integration costs and to increase flexibility in implementing the physical layer at the same time, SDR systems run functional modules—such as synchronization, modulation/de-modulation, coding/decoding, interleaving/de-interleaving, and channel parameter estimation—in a software that is executed on personal computers. Therefore, SDR technology facilitates implementation of reconfigurable radio systems where dynamic selection of parameters for the aforementioned modules is possible. The literature reports a number of SDR test-beds, designed to test network-level protocols. Popular among them are the *Wireless Open Access Research Platform* (WARP) [10, 11] developed at Rice University and the *Microsoft Research Software Radio* (MS-SORA) [12]. Hydra [13], developed at the University

of Texas at Austin, is a SDR testbed comprising radio software by GNU (recursive acronym for “GNU’s Not Unix”) and *universal software radio peripheral* (USRP) by Ettus Research™ [14]. The media access control layer (MAC) and the physical layer (PHY) design of Hydra implements the IEEE 802.11 distributed coordination function and a 2×2 MIMO-OFDM based on the IEEE 802.11a/g standard, respectively. Since then USRP hardware/GNU radio software was used to implement and test a series of heuristic receivers at PHY [15–17], other USRP/GNU radio implementations test a rate adaptation technique [18] and a random access protocol for MIMO networks [19]. To our best knowledge, there is only one publication related to a USRP hardware/GNU radio software implementation based on a theoretical framework: the expectation-maximization (EM) algorithm with a BP maximization step has been used in the context of OFDM physical-layer network coding (PNC) systems for phase tracking and single-user channel decoding [20].

This paper assesses the real performance of a MIMO-SDMA receiver, performing joint multi-user data decoding, multi-channel, and noise variance estimation (JDE), implemented on a self-made USRP/GNU radio test-bed dubbed MIMONet [21]. In contrast to many existing receiver implementations, ours is based on principled design, namely the combined BP-MF message-passing framework [7], where the virtues of BP and MF are kept but their respective drawbacks are avoided. Furthermore, the paper examines synchronization of (i) the USRP hardware, (ii) PHY burst and carrier at sample rate, and (iii) the CPU cores to process the individual data streams in parallel.

The rest of the paper is organized as follows. Section 2 presents an overview of the system under consideration, focusing on the PHY and MAC. As an example, we consider the high-rate packet-oriented IEEE 802.11n standard [22]. Section 3 presents the test-bed setup and addresses three levels of synchronization. Section 4 applies the combined BP-MF message passing to JDE. Section 5 evaluates the performance of the proposed scheme in real environments. A conventional MIMO receiver comprising a MIMO zero-forcing channel estimator and a maximum-likelihood sequence decoder is used as a benchmarking reference. We investigate the bit-error rate (BER), packet-error rate (PER), and the execution time per uncoded bit in both line-of-sight (LOS) and none-line-of-sight (NLOS) conditions. The results of our experiments clearly demonstrate the receiver complexity vs. performance trade-off: if BP-MF is executed with a single iteration, its performance is worse than that of the conventional receiver. However, the performance of the BP-MF drastically improves with the number of iterations: with five iterations, its performance is already consistently better than the conventional receiver, while up to 4 dB performance gain can

be achieved after about 20 iterations when BP-MF has converged.

Notation: In the following, $(\cdot)^\dagger$, $\Re(\cdot)$, and $\Im(\cdot)$ denote the conjugate transpose and the real and the imaginary parts of a complex argument, respectively. The symbol \angle is the argument of a complex number. The symbol $\text{diag}\{\cdot\}$ denotes a square matrix with the argument along its main diagonal. The Hadamard product of two vectors is denoted by \odot . Moreover, $\|\cdot\|$ is the two norm of the argument. Vectors and matrices in the frequency domain (time domain) are represented by boldface lowercase and uppercase Latin (Greek) letters, respectively, unless otherwise stated. The notation $\text{col}\{\cdot\}$ represents the column vector with the elements in the argument as its entries. The symbol $\mathbf{0}_{m \times n}$ denotes the $m \times n$ -dimensional all-zero matrix, whereas \mathbf{I}_n represents the identity matrix of size $n \times n$. The pdf of a multivariate Gaussian random vector with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ is denoted by $\text{CN}(\cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$. The pdf of a Gamma distribution with scale a and rate b is denoted by $\text{Ga}(\cdot; a, b)$.

2 System description

We consider a packet-oriented multi-stream MIMO-OFDM WLAN system with N_T transmit (Tx) antennas and N_R receive (Rx) antennas that implements multiple parallel, spatially segregated channels. As a working example, we refer to the IEEE 802.11n standard [22]. Each channel may support a separate data stream $k \in [1, K]$, $K \triangleq \min\{N_T, N_R\}$.

Without loss of generality, we assume that antenna $k \in [1, N_T]$ is transmitting while antenna $r \in [1, N_R]$ is receiving. The k th information stream $\{u_k[i]\}$ comprises $L_u \in \mathbb{N}$ bits per frame. The information bit sequence \mathbf{u}_k is forward error correction (FEC) encoded with rate R , fed into the symbol interleaver $\mathbf{\Pi}_k$, serial-to-parallel converted, mapped onto a modulation alphabet of size $\mathcal{S}_k = 2^{M_k}$, $M_k \in \mathbb{N}$, and the resulting symbols modulate N_a active out of N carriers. To ease channel estimation, N_p pilot tones with indices in $\mathcal{P} \subset \{(N - N_a - N_p)/2 + 1 : (N + N_a + N_p)/2 + 1\}$ are multiplexed with the data. Specifically, the pilot sequence of stream k , $k \in [1, K]$, is repeatedly taken from the k th row of the $K \times K$ dim. Walsh-Hadamard matrix, $K = 2^\ell$, $\ell \in \mathbb{N}$. Figure 1 illustrates the data and pilot placement.

The composite data vector $\mathbf{d}_k[n] \in \mathcal{S}_k^N$ of OFDM symbol n , $n \in [1, Q]$, with $Q \triangleq \lceil L_u / (RN_a) \rceil$, $L_u / R \bmod N_a = 0$, is multiplexed with the N -point Fourier matrix $\mathbf{F} \in \mathbb{C}^{N \times N}$, fed into the permutation matrix $\boldsymbol{\Psi}_s \in \mathbb{B}^{(N+P) \times N}$ to add a cyclic prefix of length P , and sent over the channel. Without loss of generality, the relative propagation delays, caused by the channel, are incorporated in the channel impulse response (CIR). In receiver r , after analog-digital conversion the equivalent time-discrete observation vector $\boldsymbol{\zeta}_r[n] \in \mathbb{C}^{N+P}$ is given by

$$\begin{aligned} \boldsymbol{\zeta}_r[n] &\triangleq \boldsymbol{\zeta}_r[nT - \tau_r] = \\ &= \mathbf{E}(e) \mathbf{T}_r(\tau_r) \sum_{k=1}^{N_T} \mathbf{A}_{rk}(\tau_r) \boldsymbol{\Psi}_s \mathbf{F}^\dagger \mathbf{d}_k[n] + \boldsymbol{\omega}_r[n], \end{aligned} \quad (1)$$

where T denotes the OFDM symbol time duration. The Toeplitz channel matrix $\mathbf{A}_{rk}(\tau_r) \in \mathbb{C}^{(N+P) \times (N+P)}$ is constituted of the static discrete-delay CIR $\boldsymbol{\alpha}_{rk} \in \mathbb{C}^L$, $L \leq P$. The distortion matrix $\mathbf{E}(e) \in \mathbb{C}^{(N+P) \times (N+P)}$ accounts for the frequency offset (FO) of the (common) local reference oscillator. It has the form

$$\mathbf{E}(e) = \text{diag}[\exp\{-j\phi\} \cdots \exp\{j\phi(N-1)\}] \quad (2)$$

with the phase $\phi \triangleq \phi(e) = 2\pi e/N$, where $e = [-0.5, 0.5]$ is the FO normalized to subcarrier spacing. The matrix $\mathbf{T}_r(\tau_r) \in \mathbb{B}^{(N+P) \times (N+P)}$, containing the time offset (TO) $\tau_r \in \mathbb{N}_0$, is given by

$$\mathbf{T}_r(\tau_r) = \begin{pmatrix} \mathbf{0}_{\tau_r \times (N+P-\tau_r)} & \mathbf{0}_{\tau_r \times \tau_r} \\ \mathbf{I}_{N+P-\tau_r} & \mathbf{0}_{(N+P-\tau_r) \times \tau_r} \end{pmatrix}. \quad (3)$$

This TO is mainly caused by the signal buffers in the receiving USRPs. Finally, the entries of the vector $\boldsymbol{\omega}_r \in \mathbb{C}^N$ are independent circularly symmetric Gaussian random variables with variance σ_w^2 .

From (1), (2), and (3), the received unsynchronized time-discrete signal vector $\boldsymbol{\zeta}_r[n]$, $r = 1, \dots, N_R$, in (1) is FO-corrected, fed into the permutation matrix $\mathbf{P}_r \in \mathbb{B}^{N \times (N+P)}$ to remove the cyclic prefix, Fourier-transformed, and finally TO-corrected by the diagonal matrix $\mathbf{D}(\tau_r) \in \mathbb{C}^{N \times N}$ with the m th diagonal entry $D_{m,m}(\hat{\tau}_r) = \exp\{j2\pi\tau_r(m-1)/N\}$, $m = 1, \dots, N$, to yield the frequency-discrete vector

$$\mathbf{y}_r[n] \triangleq \mathbf{D}(\tau_r) \mathbf{F} \mathbf{P}_r \mathbf{E}(e)^\dagger \boldsymbol{\zeta}_r[n], \quad (4)$$

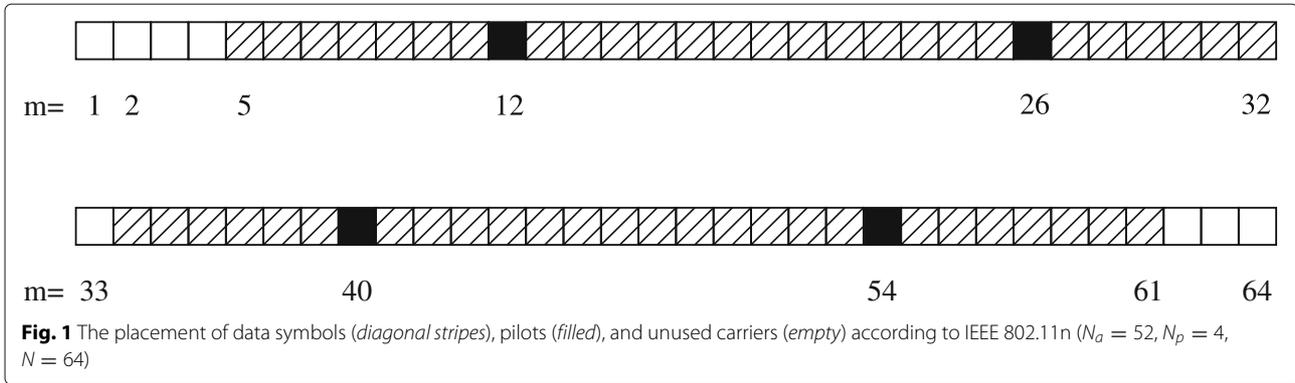
$r = 1, \dots, N_R$.

2.1 Probabilistic model of the MIMO system

Suppose that exact estimations of the TO and FO are available at the receiver. Then, the signal vector in (4) simplifies to

$$\mathbf{y}_r[n] = \sum_{k=1}^{N_T} \mathbf{h}_{rk} \odot \mathbf{d}_k[n] + \mathbf{w}_r[n], \quad (5)$$

where the components of $\mathbf{h}_{rk} \triangleq \mathbf{W} \boldsymbol{\alpha}_{rk}$, with $\mathbf{W} \in \mathbb{C}^{N \times L}$ denoting the truncated Fourier matrix, are samples of the frequency response of the channel between transmitter k and receiver r . For the design of the receiver, we make the assumption that the different channels are a priori mutually independent. Thus, we consider the prior pdf $p(\mathbf{h}_{11}, \dots, \mathbf{h}_{N_R N_T}) = \prod_{r=1}^{N_R} \prod_{k=1}^{N_T} p(\mathbf{h}_{rk})$, where each factor is modeled as a complex Gaussian prior pdf with zero mean and covariance $\boldsymbol{\Sigma}_{\mathbf{h}_{rk}}^p$, i.e.,



$$p(\mathbf{h}_{rk}) = \text{CN}(\mathbf{h}_{rk}; \mathbf{0}, \Sigma_{\mathbf{h}_{rk}}^p), \quad r \in [1, N_R], k \in [1, N_T]. \tag{6}$$

While the assumption we make does not model/take into account spatial correlation, it leads to lower complexity channel estimation [8]. We thus prefer low complexity over performance of the algorithm. In (5), the frequency-domain noise vector \mathbf{w}_r has the pdf $p(\mathbf{w}_r) = \text{CN}(\mathbf{w}_r; \mathbf{0}, \lambda_r^{-1} \mathbf{I}_N)$ with $\lambda_r \triangleq 1/N_0$ being the noise precision.

As design criterion, we use the bit-wise MAP decision rule, which minimizes the BER. That is, the decoded value of the i th information bit of the k th stream is

$$\hat{u}_k(i) = \arg \max_{u_k(i) \in \{0,1\}} p(u_k(i) | \mathcal{Y}), \tag{7}$$

where $p(u_k(i) | \mathcal{Y})$ is the marginal posterior pdf of bit $u_k(i)$ and \mathcal{Y} contains the observation vectors of all receive antennas referring to all OFDM symbols of a packet: $\mathcal{Y} \triangleq \{\mathbf{y}_r[n] \mid r \in [1, N_R], n \in [1, Q]\}$.

The marginal posterior pdf required in (7) is computed from the joint posterior pdf of the variables in the probabilistic model of the system by marginalizing out all variables but the bit of interest. Collecting all unknown system variables (channel responses, noise precision, data symbols, and information sequences) in Ψ , invoking Bayes' rule and using the system assumptions, the joint posterior pdf writes

$$p(\Psi | \mathcal{Y}) \propto \prod_{r=1}^{N_R} \left[\prod_{n=1}^Q p(\mathbf{y}_r[n] | \mathbf{h}_{r1}, \dots, \mathbf{h}_{rN_T}, \mathbf{d}_1[n], \dots, \mathbf{d}_{N_T}[n], \lambda_r) \prod_{k=1}^{N_T} p(\mathbf{h}_{rk}) \right] \times \prod_{k=1}^{N_T} p(\mathbf{d}_k[1], \dots, \mathbf{d}_k[Q] | \mathbf{u}_k) \prod_{i=1}^{I_k} p(u_k(i)). \tag{8}$$

In the following, we introduce the notation for the factors in (8) and give their functional form. We define the "observation" factors f_{Y_r} to be the likelihood of the corresponding variables. From (5), we have

$$f_{Y_r}(\mathbf{h}_{r1}, \dots, \mathbf{h}_{rN_T}, \mathbf{d}_1[1], \dots, \mathbf{d}_{N_T}[Q], \lambda_r) \triangleq \prod_{n=1}^Q p(\mathbf{y}_r[n] | \mathbf{h}_{r1}, \dots, \mathbf{h}_{rN_T}, \mathbf{d}_1[n], \dots, \mathbf{d}_{N_T}[n], \lambda_r) = \prod_{n=1}^Q \text{CN} \left(\mathbf{y}_r[n]; \sum_{k=1}^{N_T} \mathbf{h}_{rk} \odot \mathbf{d}_k[n], \lambda_r^{-1} \mathbf{I}_N \right).$$

The prior pdf of the precision of the noise at the r th receiver is denoted by f_{Λ_r} . We choose a gamma prior pdf (i.e., a conjugate prior for the likelihood f_{Y_r}) and set its parameters to be non-informative:

$$f_{\Lambda_r}(\lambda_r) \triangleq p(\lambda_r) = \text{Ga}(\lambda_r; 0, 0).$$

For the prior pdf of the channel vector \mathbf{h}_{rk} , we write $f_{H_{rk}}(\mathbf{h}_{rk}) \triangleq p(\mathbf{h}_{rk})$, with $p(\mathbf{h}_{rk})$ given by (6). The factor $p(\mathbf{d}_k[1], \dots, \mathbf{d}_k[Q] | \mathbf{u}_k)$ is denoted by f_{C_k} and stands for the deterministic operations of coding, interleaving, and modulation mapping performed in transmitter k . Finally, $f_{U_{k,i}}(u_k(i)) \triangleq p(u_k(i))$ is the prior probability mass function of the i th information bit of transmitter k . We assume a uniform prior, i.e., the bit values are a priori equally probable.

It is helpful to visualize the probabilistic dependencies between the system variables by representing the factorization of the joint posterior pdf in a factor graph [3]. With the above definitions, the factor graph representing (8) is depicted in Fig. 2.

Now that we have defined the probabilistic model, it is important to note that exact marginalization of (8) requires evaluation of high-dimensional integrals that

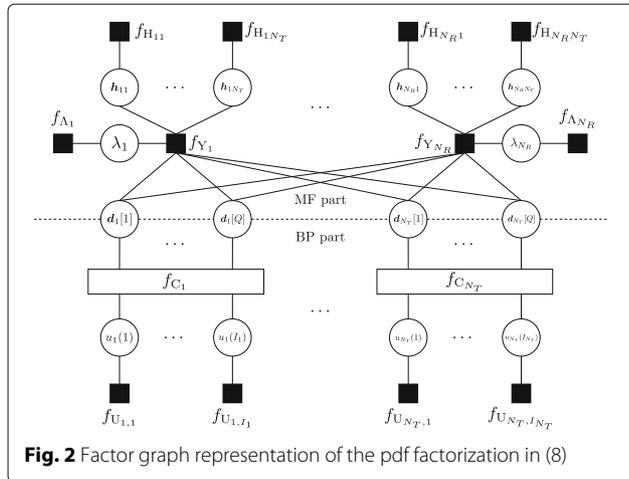


Fig. 2 Factor graph representation of the pdf factorization in (8)

do not admit closed-form expressions. That is, direct marginalization is computationally intractable. Therefore, we use an approximate inference framework to compute the estimates of the marginal pdfs of the variables, called *beliefs*. Then, the MAP decision in (7) will be applied to the beliefs $b(u_k(i)) \approx p(u_k(i)|\mathcal{Y})$ of the information bits.

2.2 Physical layer convergence protocol

To facilitate synchronization and automatic gain control, a preamble of length L_i symbols with $L_i \bmod N = 0$, is prepended to each OFDM packet.

The preamble follows the IEEE 802.11n high-throughput (HT) greenfield format without a legacy compatible part, as sketched in Fig. 3. Notice that neither coding nor scrambling is applied to generate the preamble.

The offset binary phase-shift keying (BPSK)-modulated short training field (STF) of the k th Tx antenna comprises a sequence of identical training symbols each of length $N/4$, extending over two OFDM symbols [22] Section 20.3.9. The periodic structure of the STF is ideally suited for FO estimation. A subsequent BPSK-modulated first long training field (LTF), composed of two identical training symbols each of length N , assists the receiver in estimating the TO and the CIR of the channel between the Tx antenna k and Rx antenna r . The following legacy signal

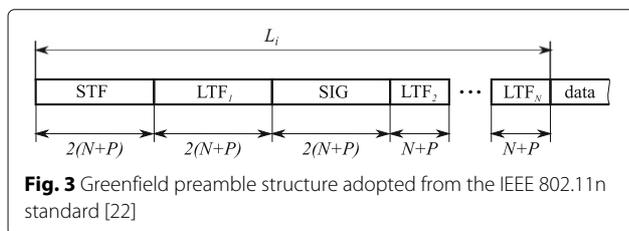


Fig. 3 Greenfield preamble structure adopted from the IEEE 802.11n standard [22]

field (SIG) carries information on the HT packet format. Additional $N_T - 1$ high-throughput LTFs $\in \mathbb{C}^{N_T(N+P)}$ are based on the same long training symbol as the first LTF in the preamble.

For the other Tx antennas $k', k' \neq k$, cyclic shift is applied to the above preamble structure to prevent beam-forming when similar signals are transmitted on different spatial streams [22].

The receiver is unsynchronized and does not know the channel coefficients and the data sequences.

3 Test-bed setup and synchronization

We consider a $N_T \times N_R$ MIMO communication link between one Tx node and one Rx node, comprising two host personal computers (PCs) and $N = N_T + N_R$ USRPs, as illustrated in Fig. 4 for $N_T = N_R = 2$.

Each node was realized with the components given in Table 1. At hardware level of SDR, each USRP contains a bank of ADC/DACs, a wideband radio front end, and a vertical antenna. At software level of SDR, digital signal processing is distributed between internal field-programmable gate arrays (FPGA) and an external host PC.

The open source software framework GNU radio under GNU general public license (GPL) was adopted to realize the transceiver chain depicted in Fig. 4. The choice of GNU radio was motivated by its scalability, its flexibility in setting the signal processing components, and its wide user base.

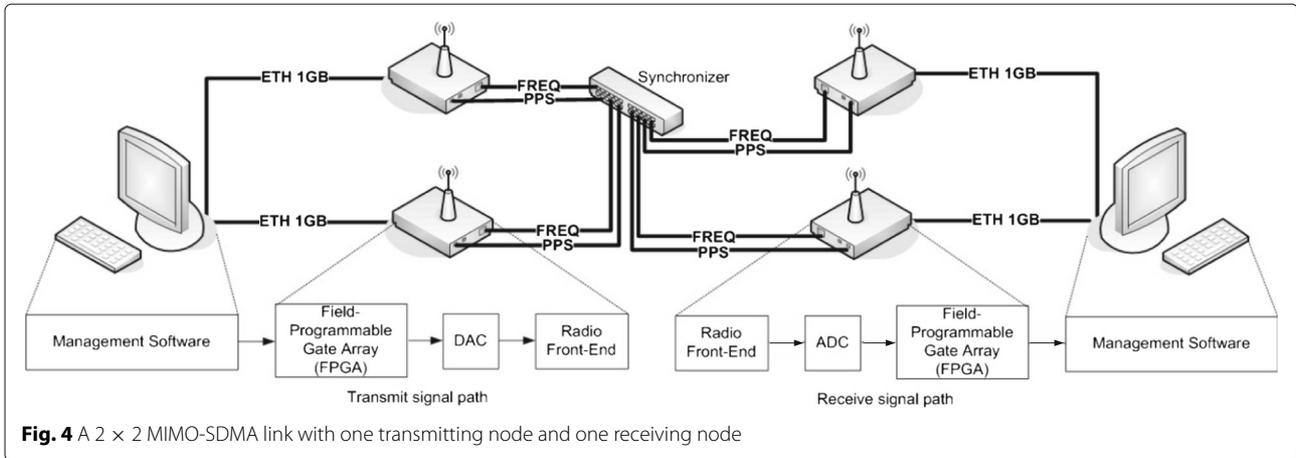
Three different levels of synchronization are needed to realize MIMO communications on computer-hosted hardware: (i) synchronization of the USRPs; (ii) burst synchronization at sample level; (iii) synchronization of the CPU cores to process the individual data streams simultaneously.

In the sequel, we describe how we addressed synchronization at each level.

3.1 USRP hardware synchronization

To enable MIMO communications, the transceiver must incorporate the following two functionalities: (i) each USRP hardware requires clock synchronization to derive the local oscillator frequency and timing synchronization to align the analog-to-digital converter/digital-to-analog converter (ADC/DAC) samples; (ii) all CPU cores have to align the digital signal streams in frequency [23] by the 10 MHz singleton and in time by the pulse-per-second (PPS) timing references. Our synchronizer, listed in Table 1, derives these references from GPS signals. In this way, PPS signals can be derived with an accuracy better than ± 50 ns.

When two USRPs are located next to each other, the internal 10 MHz /1 PPS reference of one USRP can be used to synchronize the other with a MIMO cable.



3.2 PHY joint burst and carrier synchronization

3.2.1 Methodology

Given an observation of a random vector \mathcal{Z} specified by a family $p(\mathcal{Z}|\theta)$ of pdfs parameterized by θ , the task is to compute an estimate of this parameter. The maximum likelihood (ML) point estimate of θ reads

$$\hat{\theta} = \arg \max_{\theta} p(\mathcal{Z}|\theta). \tag{9}$$

Subsequently, we present a low-complex closed-form solution to joint ML time-offset and fractional frequency-offset estimation.

3.2.2 Synchronizer design

Let us return to the specific model in (1) with $\tau_r \bmod R = 0$, $r \in [1, N_R]$, where R is the data rate. When two subsequent training symbols are identical, the received vector $\zeta[\tau_r]$ at antenna r , $r \in [1, N_R]$, and its time-shifted version $\zeta[\tau_r + (N + P)T]$ are related by

$$\zeta[\tau_r + (N + P)T] = \exp\{j2\pi e(N + P)/N\} \zeta[\tau_r]. \tag{10}$$

To obtain a ML-based synchronizer that handles frame and frequency synchronization, one might choose the parameter vector θ and observation vector \mathcal{Z} as, respectively, $\theta = \{\tau_1, \dots, \tau_{N_R}, e\}$ and $\mathcal{Z} = \text{col}\{\mathcal{Z}_1, \dots, \mathcal{Z}_{N_R}\}$ with $\mathcal{Z}_r \triangleq \zeta[\tau_r + (N + P)T] - \exp\{j2\pi e(N + P)/N\} \zeta[\tau_r]$. The latter choice accounts for the time periodicity of the preamble. The synchronizer seeks to find the joint estimate

$$\{\hat{\tau}_1, \dots, \hat{\tau}_{N_R}, \hat{e}\} = \arg \max_{\theta} \log p(\mathcal{Z}|\theta \in \mathcal{R}_{\tau,e}). \tag{11}$$

Notice that the likelihood function in (11) exhibits a unique maximum in contrast to that in the work by [24].

A valid burst synchronization is achieved if $\hat{\theta}$ is contained in the region

$$\begin{aligned} \mathcal{R}_{\tau,e} &\triangleq \left\{ \tau_1, \dots, \tau_{N_R}, e \mid \sum_{r=1}^{N_R} \zeta[\tau_r + (N + P)T] \right. \\ &= \left. \exp\{j2\pi e(N + P)/N\} \sum_{r=1}^{N_R} \zeta[\tau_r] \right\}. \end{aligned}$$

Due to the model assumptions, the log-likelihood function in (11) reads

$$\log p(\mathcal{Z}|\theta \in \mathcal{R}_{\tau,e}) \propto -\|\mathcal{Z}^\dagger \Sigma^{-1} \mathcal{Z}\|^2. \tag{12}$$

The covariance matrix Σ is given by $\Sigma = E\{\omega\omega^\dagger\} = N_0 \mathbf{I}$, where \mathbf{I} is the identity matrix. Expanding (12), we get

$$\begin{aligned} \|\mathcal{Z}^\dagger \Sigma^{-1} \mathcal{Z}\|^2 &\propto \sum_{r=1}^{N_R} \|\zeta[\tau_r]\|^2 + \|\zeta[\tau_r + (N + P)T]\|^2 \\ &\quad - 2\Re\{\exp\{j2\pi e(N + P)/N\} \zeta[\tau_r]^\dagger \\ &\quad \zeta[\tau_r + (N + P)T]\}. \end{aligned} \tag{13}$$

The log-likelihood function in (12) can first be maximized w.r.t. e , leading to

$$\hat{e} = -\frac{N}{2\pi(N + P)T} \angle \left(\sum_{r=1}^{N_R} \zeta[\tau_r]^\dagger \zeta[\tau_r + (N + P)T] \right) + I. \tag{14}$$

The integer part I of the FO e cannot be resolved at this stage. When, however, an external oscillator controls the clocks of the USRPs, $I = 0$, and e is confined to the range $[-0.5, +0.5]$. Plugging (14) into (13), the log-likelihood

Table 1 Hardware components constituting one node

Type	Model
PC	Intel®-Core™ i7-2600 @ 3.4GHz
USRPs body	Ettus Research™ N210
USRPs radio front end	Ettus Research™ XCVR2450
Synchronizer	Spectracom Corp.® EC20S

function in (12) can now be maximized w.r.t. $\tau_1, \dots, \tau_{N_R}$. Following this approach, we have

$$\{\hat{\tau}_1, \dots, \hat{\tau}_{N_R}\} = \arg \min_{\tau_1, \dots, \tau_{N_R}} \sum_{r=1}^{N_R} \left(\|\zeta[\tau_r]\|^2 + \|\zeta[\tau_r + (N+P)T]\|^2 - 2\|\zeta[\tau_r]^\dagger \zeta[\tau_r + (N+P)T]\| \right). \quad (15)$$

Inspecting (15), it can be seen that the TO of the r th receive stream is independent of the other $r-1$ receive streams. As a result,

$$\hat{\tau}_r = \arg \min_{\tau_r} \left(\|\zeta[\tau_r]\|^2 + \|\zeta[\tau_r + (N+P)T]\|^2 - 2\|\zeta[\tau_r]^\dagger \zeta[\tau_r + (N+P)T]\| \right), \quad r \in [1, N_R]. \quad (16)$$

Substituting (16) into (14), yields

$$\hat{e} = -\frac{N}{2\pi(N+P)T} \angle \left(\sum_{r=1}^{N_R} \zeta[\hat{\tau}_r]^\dagger \zeta[\hat{\tau}_r + (N+P)T] \right). \quad (17)$$

The above burst synchronization algorithm is implemented in GNU radio software at the host PC. The resulting synchronized observation vector $\mathbf{y}[n, i] \triangleq \text{col}\{y_1[n, i] \dots y_{N_R}[n, i]\}$, $n = 1, \dots, Q$, $i = 1, \dots, N$ (cf. (4)) is post-processed by MATLAB[®].

3.3 Core synchronization

GNU radio 3.6.0 has an incorporated thread-per-block scheduler that allows for each signal processing block in the flow graph to run in an independent thread. The thread, associated to one block, loops until GNU radio code is terminated. In each loop, the thread calls the block's executor. If the block has available output buffer and sufficient data in the input buffer, the executor asks for signal processing on that block and then informs neighboring block(s) about its new status. Thus, all blocks in the flow graph process incoming data chunk-by-chunk [25].

4 Iterative channel estimation and data decoding

In this section, we describe our proposed MIMO receiver algorithm, which recovers the information bits of the K data streams. The various receiver tasks—channel estimation, MIMO detection, and decoding—are jointly designed by formulating the bit recovery process as Bayesian inference on the probabilistic model of the underlying OFDM system. The resulting algorithm iteratively computes and passes messages on the factor graph representing the probabilistic model. After a fixed number of iterations (tunable parameter), the algorithm returns the most probable configuration of the bits transmitted

in the K data streams, along with estimates of other unknown quantities, such as the channel responses and the noise power.

4.1 Application to MIMO receiver design

The factor graph in Fig. 2 is split into the MF and BP parts by taking into account the functional forms of the factors and the specificities of the MF and BP algorithms. The factors f_{Y_r} , f_{λ_r} , and $f_{H_{rk}}$, $r \in [1, N_R]$, $k \in [1, N_T]$, and all variable nodes connected to them are placed in the MF part as they form a conjugate-exponential model. Given that BP has successfully been used for demapping and decoding, the rest of the factor nodes and the variable nodes connected to them represent the BP part.

4.1.1 Computation of messages and beliefs

The belief of each of the variables approximates the variable's posterior pdf. In the forthcoming computations, the following statistics will occur:

- The mean $\hat{d}_k[n, i] \triangleq \langle d_k[n, i] \rangle_{b(d_k[n, i])}$ and variance $\sigma_{d_k}^2[n, i] \triangleq \langle |d_k[n, i] - \hat{d}_k[n, i]|^2 \rangle_{b(d_k[n, i])}$ represent the (soft) estimate and uncertainty, respectively, of the symbol on the i th subcarrier of the n th OFDM symbol transmitted by k th antenna. Note that for pilot subcarriers (i.e., $i \in \mathcal{P}$), we have $\hat{d}_k[n, i] = d_k[n, i]$ and $\sigma_{d_k}^2[n, i] = 0$.
- The mean and covariance matrix of the belief $b(\mathbf{h}_{rk})$ of the respective vector of channel weights are denoted by $\hat{\mathbf{h}}_{rk} \triangleq \langle \mathbf{h}_{rk} \rangle_{b(\mathbf{h}_{rk})}$ and $\mathbf{\Sigma}_{\mathbf{h}_{rk}} \triangleq \langle (\mathbf{h}_{rk} - \hat{\mathbf{h}}_{rk})(\mathbf{h}_{rk} - \hat{\mathbf{h}}_{rk})^H \rangle_{b(\mathbf{h}_{rk})}$. The (i, i) th entry of $\mathbf{\Sigma}_{\mathbf{h}_{rk}}$ is denoted by $\sigma_{h_{rk}}^2[i]$.
- The mean $\hat{\lambda}_r \triangleq \langle \lambda_r \rangle_{b(\lambda_r)}$ represents the estimate of the precision of the r th receiver's noise.

4.1.1.1 Channel estimation

Obtaining the beliefs $b(\mathbf{h}_{rk}) \approx p(\mathbf{h}_{rk}|\mathcal{Y})$, $r \in [1, N_R]$, $k \in [1, N_T]$ corresponds to channel estimation and requires computing the MF messages related to the channel vectors. We readily show that the message from the observation node f_{Y_r} has the Gaussian form

$$m_{f_{Y_r} \rightarrow \mathbf{h}_{rk}}^{\text{MF}}(\mathbf{h}_{rk}) \propto \prod_{i \in \mathcal{D} \cup \mathcal{P}} \prod_{n=1}^Q \text{CN} \left(\mathbf{h}_{rk}[i]; \tilde{h}_{rk}^{\text{obs}}[n, i], \sigma_{h_{rk}}^2[n, i] \right),$$

where

$$\tilde{h}_{rk}^{\text{obs}}[n, i] = \frac{\hat{d}_k^*[n, i]}{\sigma_{d_k}^2[n, i] + |\hat{d}_k[n, i]|^2} \left(y_r[n, i] - \sum_{k' \neq k} \hat{h}_{rk'}[i] \hat{d}_{k'}[n, i] \right)$$

$$\sigma_{h_{rk}}^2[n, i] = \hat{\lambda}_r^{-1} \left(\sigma_{d_k}^2[n, i] + |\hat{d}_k[n, i]|^2 \right)^{-1}$$

represents the estimates and their uncertainty when taking into account only the observations from the n th

OFDM symbols and no prior information. Using the fact that multiplying Gaussian pdfs results in a Gaussian pdf, we obtain

$$m_{f_{Y_r} \rightarrow \mathbf{h}_{rk}}^{\text{MF}}(\mathbf{h}_{rk}) \propto \prod_{i \in \mathcal{D} \cup \mathcal{P}} \text{CN}\left(\mathbf{h}_{rk}[i]; \hat{\mathbf{h}}_{rk}^{\text{obs}}[i], \sigma_{\hat{\mathbf{h}}_{rk}^{\text{obs}}[i]}^2\right)$$

with

$$\sigma_{\hat{\mathbf{h}}_{rk}^{\text{obs}}[i]}^2 = \left(\sum_{n=1}^Q \frac{1}{\sigma_{\hat{\mathbf{h}}_{rk}^{\text{obs}}[n,i]}^2} \right)^{-1}$$

$$\hat{\mathbf{h}}_{rk}^{\text{obs}}[i] = \sigma_{\hat{\mathbf{h}}_{rk}^{\text{obs}}[i]}^2 \sum_{n=1}^Q \frac{\tilde{\mathbf{h}}_{rk}^{\text{obs}}[n,i]}{\sigma_{\hat{\mathbf{h}}_{rk}^{\text{obs}}[n,i]}^2}.$$

Defining the vector $\hat{\mathbf{h}}_{rk}^{\text{obs}} \triangleq \left(\hat{\mathbf{h}}_{rk}^{\text{obs}}(i) \mid i \in \mathcal{D} \cup \mathcal{P} \right)^T$ and the diagonal matrix $\Sigma_{\mathbf{h}_{rk}}^{\text{obs}}$ whose (i, i) th entry equals $\sigma_{\hat{\mathbf{h}}_{rk}^{\text{obs}}[i]}^2$, we write $m_{f_{Y_r} \rightarrow \mathbf{h}_{rk}}^{\text{MF}}(\mathbf{h}_{rk}) \propto \text{CN}\left(\mathbf{h}_{rk}; \hat{\mathbf{h}}_{rk}^{\text{obs}}, \Sigma_{\mathbf{h}_{rk}}^{\text{obs}}\right)$. Given that $m_{f_{\mathbf{h}_{rk}} \rightarrow \mathbf{h}_{rk}}^{\text{MF}}(\mathbf{h}_{rk}) = f_{\mathbf{H}_{rk}}(\mathbf{h}_{rk})$, we obtain that the beliefs of \mathbf{h}_{rk} , $r \in [1, N_R]$, $k \in [1, N_T]$, are Gaussian pdfs:

$$b(\mathbf{h}_{rk}) \propto m_{f_{\mathbf{h}_{rk}} \rightarrow \mathbf{h}_{rk}}^{\text{MF}}(\mathbf{h}_{rk}) m_{f_{Y_r} \rightarrow \mathbf{h}_{rk}}^{\text{MF}}(\mathbf{h}_{rk})$$

$$= \text{CN}\left(\mathbf{h}_{rk}; \hat{\mathbf{h}}_{rk}, \Sigma_{\mathbf{h}_{rk}}\right) \quad (18)$$

with mean and covariance matrix

$$\Sigma_{\mathbf{h}_{rk}} = \left(\Sigma_{\mathbf{h}_{rk}}^{\text{P}} \right)^{-1} + \left(\Sigma_{\mathbf{h}_{rk}}^{\text{obs}} \right)^{-1}$$

$$\hat{\mathbf{h}}_{rk} = \Sigma_{\mathbf{h}_{rk}} \left(\Sigma_{\mathbf{h}_{rk}}^{\text{obs}} \right)^{-1} \hat{\mathbf{h}}_{rk}^{\text{obs}}.$$

4.1.1.2 Estimation of the noise precision

The message from the observation factor node $m_{f_{Y_r} \rightarrow \lambda_r}^{\text{MF}}(\lambda_r)$ is found to be proportional to a gamma pdf. Given that the prior is a non-informative gamma pdf, the belief of the noise precision at receiver r equals the message $m_{f_{Y_r} \rightarrow \lambda_r}^{\text{MF}}(\lambda_r)$. Therefore,

$$b(\lambda_r) = \text{Ga}(\lambda_r; QL_d + 1, \beta_r), \quad r \in [1, N_R], \quad (19)$$

where the rate of the gamma pdf is given by

$$\beta_r = \sum_{i \in \mathcal{D} \cup \mathcal{P}} \sum_{n=1}^Q \left[\left| y_r[n, i] - \sum_k \hat{\mathbf{h}}_{rk}[i] \hat{d}_k[n, i] \right|^2 + \sum_k \sigma_{\hat{d}_k}^2[n, i] \sigma_{\hat{\mathbf{h}}_{rk}}^2[i] + \sum_k \sigma_{\hat{\mathbf{h}}_{rk}}^2[i] \left| \hat{d}_k[n, i] \right|^2 + \sum_k \sigma_{\hat{d}_k}^2[n, i] \left| \hat{\mathbf{h}}_{rk}[i] \right|^2 \right].$$

The estimates of the noise precisions are $\hat{\lambda}_r = \frac{QL_d}{\beta_r}$, $r \in [1, N_R]$.

4.1.1.3 MIMO detection and decoding

The messages from the observation nodes are found to be

$$m_{f_{Y_r} \rightarrow \mathbf{d}_k[n]}^{\text{MF}}(\mathbf{d}_k[n]) \propto \prod_{i \in \mathcal{D}} \text{CN}\left(\mathbf{d}_k[n, i]; \hat{d}_k^{\text{obs}}[n, i], \sigma_{\hat{d}_k^{\text{obs}}[n, i]}^2\right), \quad (20)$$

$n \in [1, Q]$, $k \in [1, N_T]$, where

$$\hat{d}_k^{\text{obs}}[n, i] = \frac{\hat{h}_{rk}^*[i]}{\sigma_{\hat{\mathbf{h}}_{rk}}^2[i] + |\hat{\mathbf{h}}_{rk}[i]|^2} \left(y_r[n, i] - \sum_{k' \neq k} \hat{\mathbf{h}}_{rk'}[i] \hat{d}_{k'}[n, i] \right)$$

$$\sigma_{\hat{d}_k^{\text{obs}}[n, i]}^2 = \hat{\lambda}_r^{-1} \left(\sigma_{\hat{\mathbf{h}}_{rk}}^2[i] + |\hat{\mathbf{h}}_{rk}[i]|^2 \right)^{-1}.$$

Note that the right-hand side of (20) is evaluated at the symbol constellation points. When normalized, those discrete messages “carry” extrinsic information on the different constellation points.

For all $n \in [1, Q]$, the messages $\prod_{r=1}^{N_R} m_{f_{Y_r} \rightarrow \mathbf{d}_k[n]}^{\text{MF}}(\mathbf{d}_k[n])$ are passed to the BP part. They represent the input to the demappers and decoders, which compute messages related to the coded bits and information bits using BP. For example, applying BP to decode convolutional codes is equivalent to using the BCJR algorithm [26]. The messages $m_{f_{C_k} \rightarrow \mathbf{d}_k[n]}^{\text{BP}}(\mathbf{d}_k[n])$, $k \in [1, N_T]$, $n \in [1, Q]$ contain extrinsic information on the respective symbols obtained from the decoders and soft mappers. The symbol beliefs

$$b(\mathbf{d}_k[n]) \propto \prod_{r=1}^{N_R} m_{f_{Y_r} \rightarrow \mathbf{d}_k[n]}^{\text{MF}}(\mathbf{d}_k[n]) m_{f_{C_k} \rightarrow \mathbf{d}_k[n]}^{\text{BP}}(\mathbf{d}_k[n]), \quad (21)$$

$k \in [1, N_T]$, $n \in [1, Q]$, represent approximations of the a posteriori probabilities (APPs) of the symbols. These values are further passed to the MF part.

The decoders also output the messages $m_{f_{C_k} \rightarrow u_k[i]}^{\text{BP}}(u_k[i])$, $k \in [1, N_T]$, $i \in [1, I_{N_k}]$. Given the prior pdfs, the beliefs of the information bits are obtained as

$$b(u_k[i]) \propto m_{f_{C_k} \rightarrow u_k[i]}^{\text{BP}}(u_k[i]) f_{U_{k,i}}(u_k(i)).$$

4.1.2 Outline of the iterative algorithm

We now define the iterative algorithm by specifying a schedule for the message computations.

The BP-MF algorithm needs to be initialized. First, the algorithm sets the conditional expectations $\hat{\lambda}_r = QL_d / \sum_{n=1}^Q \|\mathbf{y}_r[n]\|^2$ and $\hat{\mathbf{h}}_{rk} = \mathbf{0}$, $k \in [1, N_T]$, $r \in [1, N_R]$. Then, the beliefs of those subvectors of \mathbf{h}_{rk} , corresponding to the pilot indices \mathcal{P} , are computed successively for each k . Next, $\hat{\mathbf{h}}_{rk}^{\text{obs}}[i] = \mathbf{0}$ and $\sigma_{\hat{\mathbf{h}}_{rk}^{\text{obs}}[i]}^2 = \mathbf{0}$, for all $i \in \mathcal{D}$, and the beliefs of the channel vectors are computed as \mathbf{h}_{rk} , $r \in [1, N_R]$, $k \in [1, N_T]$. Having obtained the initial estimates of the beliefs of the channel weights and noise precision, the algorithm then performs MIMO detection

with the initial parameter setting $\hat{d}_k[n, i] = 0$ and variance $\sigma_{d_k}^2[n, i] = 1, k \in [1, N_T], n \in [1, Q], i \in [1, I_{N_k}]$. The beliefs $b(\hat{d}_k[n])$, $n \in [1, Q]$ are computed sequentially for each $k \in [1, N_T]$, a scheme which resembles successive interference cancellation. The initial stage ends with demapping and decoding.

During subsequent iterations, messages for soft mapping, channel and noise precision estimation, MIMO detection, and demapping and decoding are computed.

After convergence, the information bits are determined by taking hard decisions based on the beliefs $b(u_k[i]), k \in [1, N_T], i \in [1, I_{N_k}]$.

5 Performance evaluation in real environments

The performance of the BP-MF iterative receiver was experimentally evaluated at some premises of the Istituto di Informatica e Telematica (IIT) at the National Research Council (CNR), Pisa, Italy. The building is constructed of concrete with steel reinforcement and with wooden doors. The measurements were conducted in the first floor in the Algorithms and Computational Mathematics laboratory of the institute. A map of this area is given in Fig. 5. Test cases include (i) a LOS link in an office environment, comprising the windows, office furniture, and computers and (ii) a non-LOS link, striking office rooms, bathrooms, and a corridor. Communication is based on the IEEE 802.11n standard for OFDM-MIMO [22]. The settings of the OFDM-SDMA system emulated on the test-bed are outlined in Table 2. The number of

counted error events per simulation point is large enough to produce sufficiently tight confidence intervals. Hence, the confidence intervals are omitted in subsequent plots.

The BP-MF receiver was benchmarked against a (low-complex) conventional MIMO receiver, composed of a linear MIMO channel filter using pilot-based channel estimates and a bank of individual ML-sequence decoders. Specifically, the MIMO channel estimate is based on the least squares technique given the Walsh-Hadamard pilot matrix and the synchronized observation matrix \mathbf{y} . The composite MIMO channel estimates at all active tones, $\hat{\mathbf{H}}$, is obtained by piecewise linear interpolation. The decorrelating MIMO multiuser detector outputs the signal $\hat{\mathbf{x}} \triangleq \hat{\mathbf{H}}^{-1} \mathbf{y}$.

In real environments, the SNR at the individual slicer is unknown. To obtain a guess of the SNR, though, we use the following histogram techniques. Let

$$\hat{\gamma}_{b,k}[n] \triangleq \frac{\hat{\sigma}_b^2}{R \hat{\sigma}_{w,k}^2[n]} \tag{22}$$

be the packet-SNR estimate per uncoded bit of OFDM symbol $n \in [1, Q]$ with signal power $\hat{\sigma}_b^2 = \hat{\mathbf{H}}^{-1} \mathbf{H} \triangleq 1$ and packet noise power $\hat{\sigma}_{w,k}^2[n]$. For known transmission symbols, the latter quantity is Gamma-distributed with conditional expectation

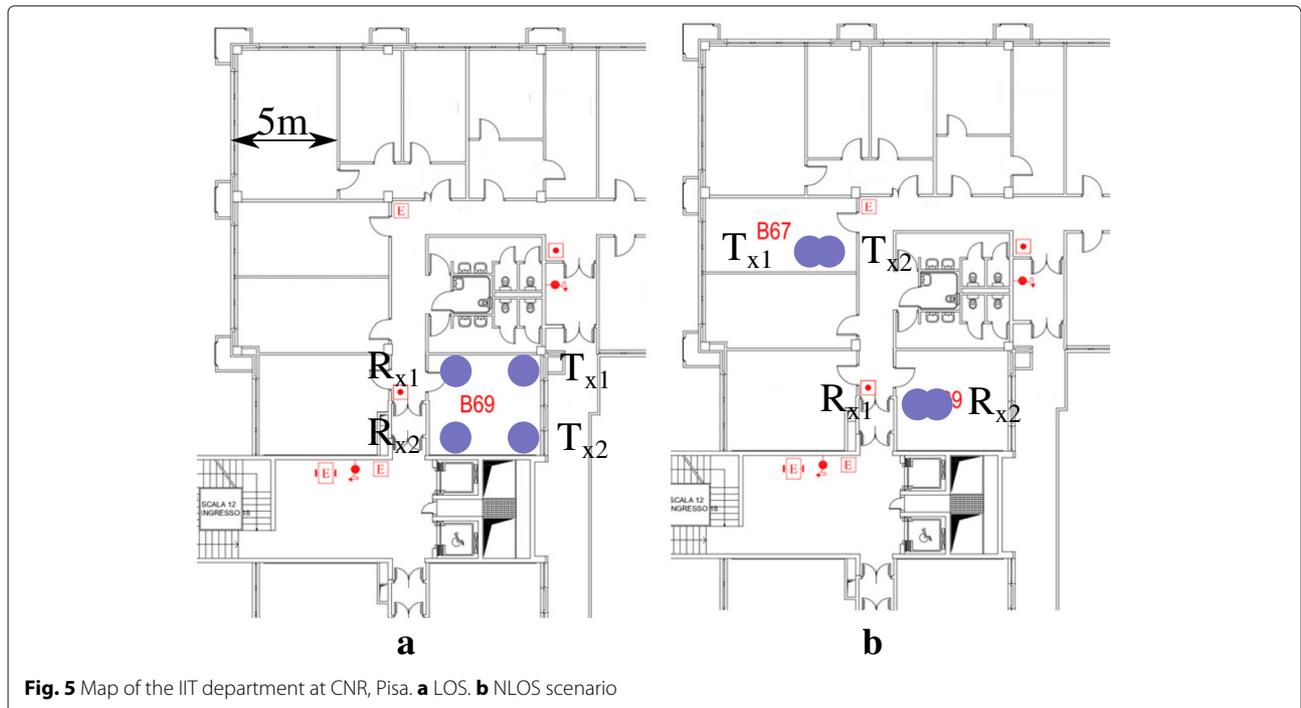


Fig. 5 Map of the IIT department at CNR, Pisa. a LOS. b NLOS scenario

Table 2 Parameter settings of the considered OFDM-SDMA system

Parameters	Value
Number of streams	$K = 2$
Number of antennas	$N_T = 2, N_R = 2$
Tx antenna gain	$G_k = \{10, 20, 30\}$ dB
Rx antenna gain	$G_r = \{10, 20, \dots, 60\}$ dB
Center frequency	2.4 GHz
Sampling frequency	200 kHz
FFT size	$N = 64$
Number of active subcarriers	$N_a = 52$
Number of pilot symbols	$N_p = 4$
FEC	$R = 1/2, G = (171, 133)_8$
Number of preamble OFDM symbols	$Q = 7$
Modulation	BPSK
Payload size	$L_u = \{96, 512, 1032\}$ bytes
Number of transmitted packets for gain tuple $\{G_k, G_r\}$	2000

$$\hat{\sigma}_{w,k}^2[n] \triangleq E[|w_k|^2 | \mathbf{x}_k, d_k] = \frac{\sum_{i \in \mathcal{D}} |x_k[n, i] - d_k[n, i]|^2}{N_a}.$$

To construct a histogram from $\hat{\gamma}_{b,k}$, the data is split into bins of width 0.5 dB. Each bin contains the number of occurrences of scores in the data set $\hat{\gamma}_{b,k}[n]$, $n \in [1, Q]$, which fall within that bin. For the sake of fair comparison, both conventional MIMO and BP-MF receivers compute the SNR in the same way.

Experiments were conducted during nights or weekends, to avoid as much as possible dynamic interference with employees' movements and devices.

5.1 NLOS scenario

We first present the experimental results obtained in NLOS conditions. Without loss of generality, we focus on stream 1. Surprisingly, the BP-MF receiver was initially not able to decode the individual data streams at the output of the MIMO synchronizer (4) while it had done so in synthetic channels [8]. Closer examination revealed that the real channel is strongly correlated in space, a property that had not been accounted in the underlying channel model of the BP-MF receiver. When it is included in the system model, spatial channel correlation can be jointly estimated and exploited to improve the accuracy of the estimates of the other system parameters and vice versa [8]. As this estimation problem is outside the scope of the paper, spatial correlation is subsequently suppressed by an equalizer with transfer function $\hat{\mathbf{H}}^{-1}$ prior to MIMO reception. The same approach is pursued in the conventional MIMO receiver.

Figure 6 shows the bit-error rate (BER) vs. quantized packet-SNR per bit with the number of iterations i as a parameter. The payload size is $L_u = 512$ bytes. Inspecting Fig. 6, we see that with increasing iteration index i , the BER performance keeps improving until a minimum of the variational free energy is found. At the low and high SNR regimes, the BP-MF algorithm converges to a fix-point around $i \triangleq i^* = 17$ iterations. After convergence, the proposed BP-MF outperforms the conventional receiver over the entire SNR range with a gain of up to 4 dB. This gain comes despite the mismatch between the real propagation conditions and those mimicked by the model used to derive the BP-MF algorithm.

The packet-error rate (PER) vs. $[\hat{\gamma}_b]$ with the number of iterations as a parameter is reported in Fig. 7 for a packet length of $L_u = 512$ bytes. To achieve a typical target PER of $\text{PER} = 0.1$, the BP-MF algorithm requires $[\hat{\gamma}_b] \approx 2.1$ dB after convergence. The conventional MIMO receiver, in contrast, meets the target-PER only at $[\hat{\gamma}_b] \approx 5.4$ dB, leaving a gap of approximately 3.3 dB. The performance gap between the BP-MF and conventional MIMO receivers can be as high as 4 dB depending on the target PER value.

With increasing payload size L_u , the BP-MF receiver shows improved PER performance, as shown in Fig. 8. This is mainly because the individual bit errors are independent so that their impact on the packet-error rate is only $1 - (1 - \text{BER})^{L_u}$. The target packet error-rate $\text{PER} = 0.1$ is met at $[\hat{\gamma}_b] \approx 1.9$ (4.1) dB for $L_u = 96$ ($L_u = 1032$) bytes, corresponding to a SNR gap of 2.1 dB for a tenfold increase in payload size. Also shown in Fig. 8, the (one-stage) conventional receiver, in contrast, has difficulties in handling large packet sizes.

5.2 LOS scenario

We now discuss the experimental results obtained in LOS conditions. From Fig. 9, it can be seen that the BER curve of the BP-MF receiver is confined in a narrow range around $\text{BER} = 0.2$ for small SNR values up to $[\hat{\gamma}_b] = 2$ dB, followed by a waterfall region for higher SNR values. Convergence is achieved after approximately $i^* = 10$ iterations. Generally speaking, the performance of the BP-MF receiver in the LOS condition is inferior to that in the NLOS condition, mainly because the number of degrees of freedom of the channel that can be exploited by the receiver is higher in the latter condition. The BP-MF receiver converges faster in LOS than it does in NLOS, because its convergence speed is inversely proportional to the number of degrees of freedom. The BP-MF receiver performs roughly 2 dB better than the conventional MIMO receiver over the entire SNR range.

The behavior of the PER curve in Fig. 10 is similar to that of the BER curve in Fig. 9. The waterfall region starts around $[\hat{\gamma}_b] = 2$ dB. The target $\text{PER} = 0.1$ is met at $[\hat{\gamma}_b] = 3.5$ dB, i.e., 1.4 dB higher than in NLOS. Notice that

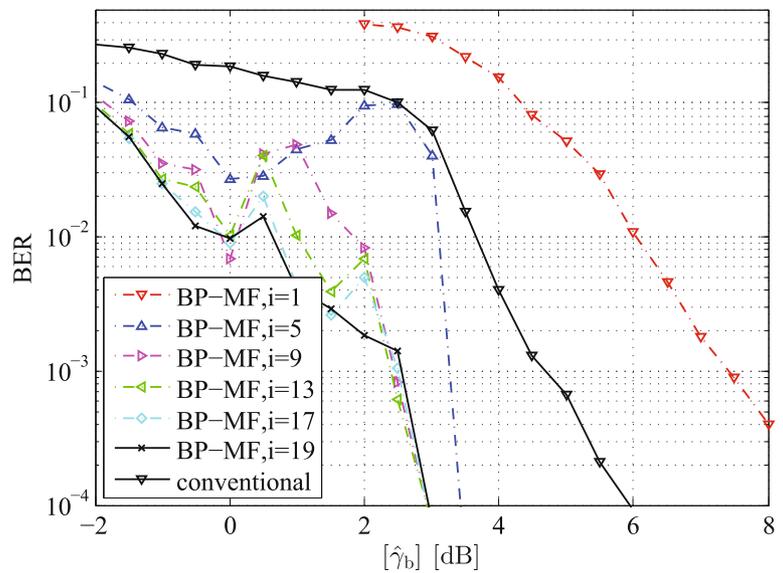


Fig. 6 BER performance of the BP-MF receiver in a NLOS condition for different values of the number of iterations ($L_u = 512$ bytes)

the PER curve of the conventional MIMO receiver only starts decreasing towards zero beyond the SNR range of interest.

With increasing payload size, the PER performance of the BP-MF receiver still improves. However, the target of $PER = 0.1$ is met at $[\hat{\gamma}_b] \approx 3.1$ (6.8) dB for $L_u = 96$ ($L_u = 1032$) bytes. The gap in SNR is now 3.7 dB (Fig. 11).

Finally, Fig. 12 shows the execution time per uncoded bit of the BP-MF receiver normalized to the execution time by the conventional receiver. Payload sizes of $L_u = 96$ and

$L_u = 1032$ bytes are considered. As the BP-MF receiver operates at symbol level, its execution time is roughly independent of the payload size and, therefore, we conclude that the per-bit computational complexity is linear in the payload size plus some offset. After the first iteration, the execution of the BP-MF algorithm has already required ten times more time than the execution of the conventional MIMO receiver. Ergo, the offset is ten. After convergence, typically at $i^* = 17$ iterations, the ratio is as high as 87.

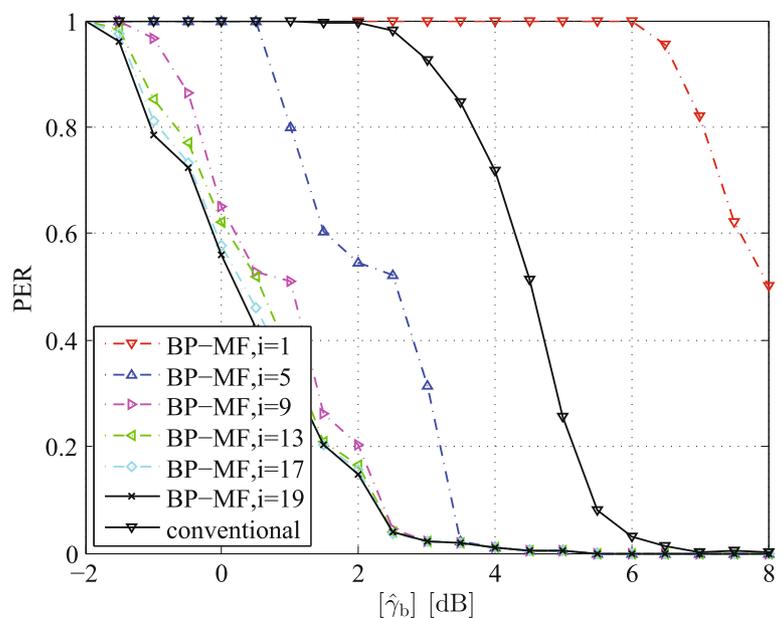


Fig. 7 PER performance of the BP-MF receiver in a NLOS condition for different values of the number of iterations ($L_u = 512$ bytes)

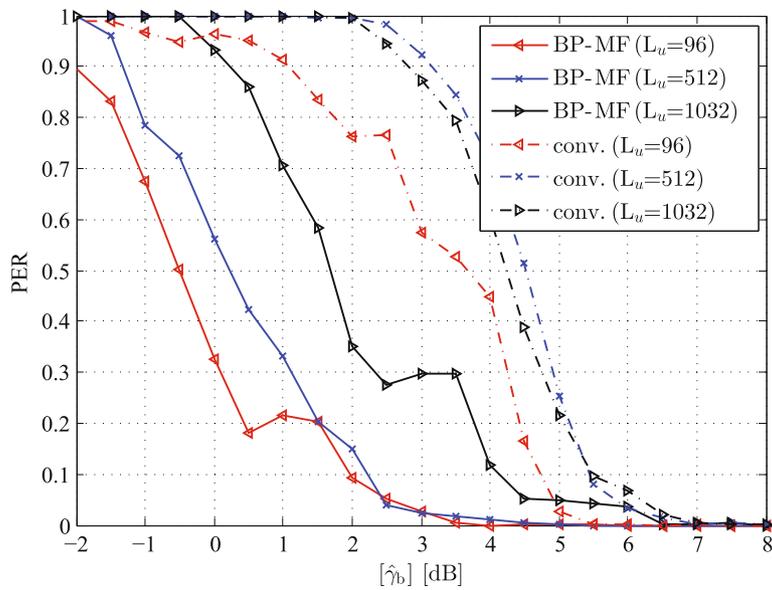


Fig. 8 PER performance of the BP-MF receiver in a NLOS condition for different payload sizes ($i^* = 17$ iterations)

6 Discussion and conclusions

In this paper, we investigated for the first time the trade-off between complexity, running time, and performance for an advanced, iterative MIMO-SDMA receiver operating in real world conditions. The receiver was derived within a unified message-passing framework that combines belief propagation and mean-field approximation. At each iteration, messages related to the channel parameters and noise precision are passed from the mean

field part to the belief propagation part of the factor graph for the belief of the data and vice versa. The latter part represents the probabilistic model of the communication system. The proposed receiver was implemented in USRP/GNU radio.

The study showed that while substantial performance improvements with respect to a conventional receiver can be achieved—ranging from 2 to 4 dB depending on packet size and LOS conditions—these benefits come

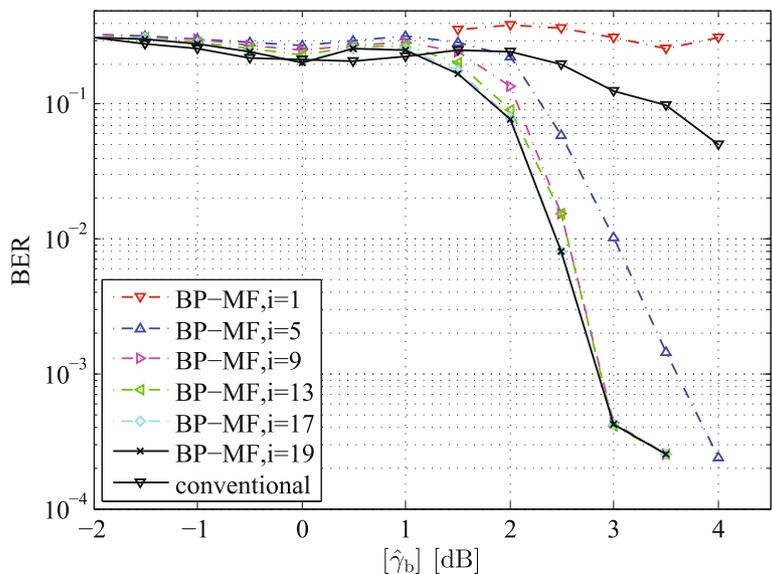


Fig. 9 BER performance of the BP-MF receiver in a quasi LOS condition for different values of the number of iterations ($L_u = 512$ bytes)

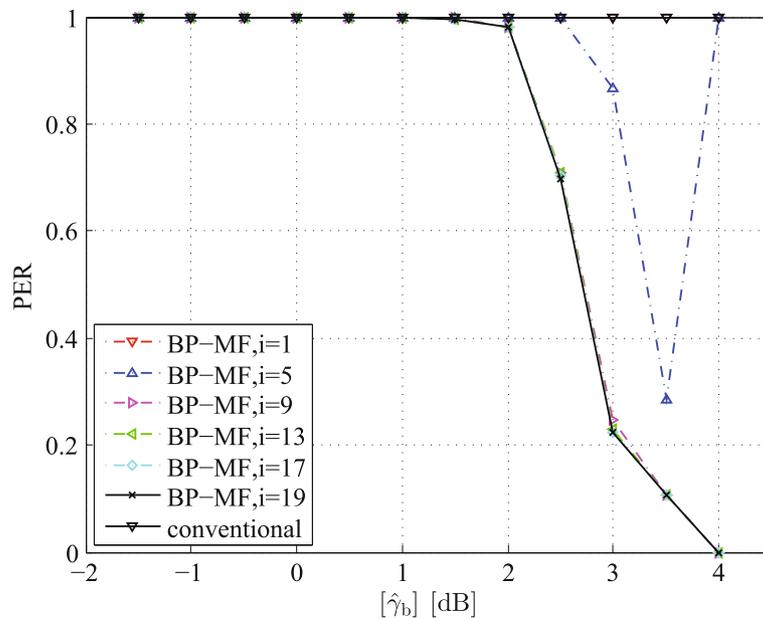


Fig. 10 PER performance of the BP-MF receiver in a quasi-LOS condition with the number of iterations as parameter ($L_u = 512$ bytes)

at an increase in per-bit decoding running time that increases linearly with the number of performed iterations. If full convergence (i.e., best performance) is sought, decoding running time can be as much as 2 orders of magnitude larger than that of a conventional receiver. However, substantial performance improvements can be

achieved also with a smaller number of iterations, especially in a rich scattering environment (NLOS scenario). Summarizing, our study clearly shows that in practical settings the trade-off between receiver complexity, running time, and performance should be carefully evaluated to strike the best compromise between these metrics.

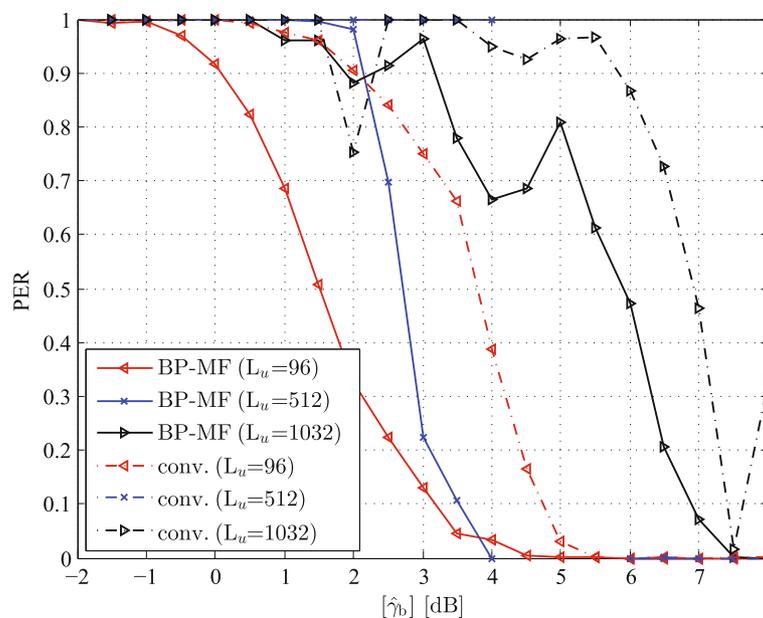
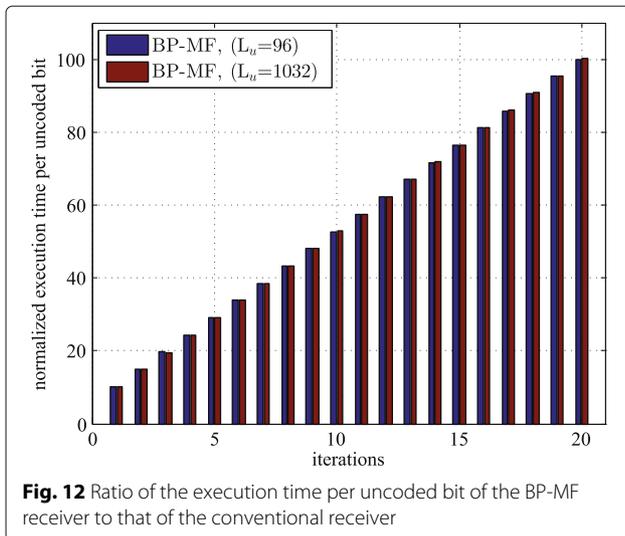


Fig. 11 PER performance of the BP-MF receiver in a quasi-LOS condition for different payload sizes ($i = 10$ iterations)



The results presented in this paper should be considered as a first step towards gaining an understanding of the feasibility of deploying advanced, iterative MIMO-SDMA receivers performance in real-world conditions. Future work includes considering higher-order modulations, as well as more complex MIMO configurations, including distributed MIMO channels.

Authors' contributions

The paper presents a novel software radio implementation of a MIMO-SDMA receiver based on a recently developed unified message-passing framework. All authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Author details

¹Department of Computer Science, University of Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy. ²Department of Electronic Systems, Aalborg University, Frederik Bajers Vej 7, 9220 Aalborg, Denmark. ³Istituto di Informatica e Telematica, National Research Council, Via G. Moruzzi 1, 56124 Pisa, Italy. ⁴SENSEable City Lab, Massachusetts Institute of Technology, 77 Massachusetts Avenue, 02139 Cambridge, MA USA.

Received: 11 September 2015 Accepted: 6 December 2016

Published online: 03 January 2017

References

1. CM Bishop, *Pattern Recognition and Machine Learning*. (Springer, Secaucus, 2006)
2. J Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. (Morgan Kaufmann Publishers Inc., San Francisco, 1988)
3. FR Kschischang, BJ Frey, H-A Loeliger, Factor graphs and the sum-product algorithm. *IEEE Trans. Inform. Theory*. **47**(2), 498–519 (2001)
4. H-A Loeliger, J Dauwels, J Hu, S Korl, L Ping, FR Kschischang, The factor graph approach to model-based signal processing. *Proc. IEEE*. **95**(6), 1295–1322 (2007). doi:10.1109/JPROC.2007.896497
5. M-A Badiu, GE Kirkelund, C Navarro Manchón, E Riegler, BH Fleury, in *Proc. IEEE Int. Symp. Inf. Th. (ISIT 2012)*. Message-passing algorithms for channel estimation and decoding using approximate inference, (Cambridge, 2012), pp. 2386–2390
6. J Winn, CM Bishop, Variational message passing. *J. Mach. Learn. Res.* **6**, 661–694 (2005)

7. E Riegler, GE Kirkelund, C Navarro Manchón, M-A Badiu, BH Fleury, Merging belief propagation and the mean field approximation: a free energy approach. *IEEE Trans. Inform. Theory*. **59**(1), 588–602 (2013)
8. C Navarro Manchón, GE Kirkelund, E Riegler, L Christensen, BH Fleury, Receiver architectures for MIMO-OFDM based on a combined VMP-SP algorithm (2011). arXiv:1111.5848 [stat.ML]. <http://arxiv.org/pdf/1111.5848.pdf>
9. M Dillinger, K Madani, N Alonistioti, *Software Defined Radio: Architectures, Systems and Functions*. (John Wiley & Sons, NJ, 2003)
10. H Yu, L Zhong, A Subharwal, D Kao, in *Proc. ACM Mobicom*. Beamforming on mobile devices: a first study, (Las Vegas, 2011)
11. E Aryafar, N Anand, T Salonidis, E Knightly, in *Proc. ACM Mobicom*. Design and experimental evaluation of multi-user beamforming in wireless lans, (Chicago, 2010)
12. K Tan, H Liu, J Fang, W Wang, J Zhang, M Chen, GM Voelker, in *Proc. ACM Mobicom*. SAM: enabling practical spatial multiple access in wireless LAN, (Beijing, 2009)
13. K Mandke, S-H Choi, G Kim, R Grant, RC Daniels, W Kim, RWJ Heath, SM Nettles, in *Proc. IEEE 65th Vehicular Technology Conference (VTC2007-Spring)*. Early results on Hydra: A flexible MAC/PHY multihop testbed, (Dublin, 2007), pp. 1896–1900
14. EttusResearch: <http://www.ettus.com/>. Accessed 14 Dec 2016
15. S Gollakota, SD Perli, D Katabi, in *Proc. ACM SIGCOMM*. Interference alignment and cancellation, (Barcelona, 2009), pp. 159–170
16. P Zetterberg, NN Moghadam, in *Systems, Signals and Image Processing (IWSSIP), 2012 19th International Conference On*. An experimental investigation of SIMO, MIMO, interference-alignment (IA) and coordinated multi-point (CoMP), (Vienna, 2012), pp. 211–216
17. P Eliardsson, U Uppman, in *Proc. 7th Karlsruhe Workshop on Software Radios (WSR'12)*. An SDR implementation of a MIMO communication testbed, (Karlsruhe, 2012)
18. WL Shen, YC Tung, KC Lee, KC Lin, S Gollakota, D Katabi, MS Chen, in *Proc. ACM Mobicom*. Rate adaptation for 802.11 multiuser MIMO networks, (Istanbul, 2012)
19. KC Lin, S Gollakota, D Katabi, in *Proc. ACM SIGCOMM*. Random access heterogeneous MIMO networks, (Toronto, 2011), pp. 146–157
20. T Wang, SC Liew, L You, in *Proc. of the 2014 ACM Workshop on Software Radio Implementation Forum*. SRIF '14. Joint phase tracking and channel decoding for OFDM PNC: algorithm and experimental evaluation (ACM, New York, 2014), pp. 69–76. doi:10.1145/2627788.2627792
21. F Martelli, A Kocian, P Santi, V Gardellin, in *Proceedings of the 2014 ACM Software Radio Implementation Forum*. SRIF '14. MIMO-OFDM spatial multiplexing technique implementation for GNU radio (ACM, Chicago, 2014), pp. 85–92. doi:10.1145/2627788.2627795
22. IEEE 802.11-2012 (Clause 20), Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications (2012). doi:10.1109/IEEESTD.2012.6178212
23. EttusResearch: Application note on synchronization and MIMO capability with USRP devices. http://www.ettus.com/content/files/kb/mimo_and_sync_with_usrp_updated.pdf. Accessed 14 Dec 2016
24. A Saemi, V Meghadi, P-J Cances, MR Zahabi, Joint ML time-frequency synchronization and channel estimation algorithm for MIMO-OFDM systems. *IET Circuits, Devices and Systems*, 103–111 (2008). doi:10.1049/iet-cds:20070024
25. F Ge, C-YJ Chiang, Y Gottlieb, R Chadha, in *Proc. IEEE Global Telecommunications Conference 2011 (GLOBECOM 2011)*. GNU radio-based digital communications: computational analysis of a GMSK transceiver, (Houston, 2011), pp. 1–6
26. L Bahl, J Cocke, F Jelinek, J Raviv, Optimal decoding of linear codes for minimizing symbol error rate. *IEEE Trans. Inform. Theory*. **20**(3), 284–287 (1974)